The International Archives of the Photogrammetry, Remote Sensing and Spatial Information Sciences, Volume L-4/W3-2026
10th International Conference on Smart Data and Smart Cities (SDSC), 29 September–1 October 2026, Sofia, Bulgaria

From Building Archetypes to Data-Driven Definitions: The Impact of Window-to-Wall Granularity on Urban Building Energy Modelling

Camilo León-Sánchez[1], Taşkın Özkan[2], Giorgio Agugiaro[2]

[1] Laboratory of Geo-information Science and Remote Sensing, Wageningen University & Research, Gelderland, Wageningen, 6700 AA, The Netherlands – camilo.leonsanchez@wur.nl

[2] 3D Geoinformation Group, Department of Urbanism, Faculty of Architecture and Built Environment, Delft University of Technology, 2628 BL Delft, The Netherlands – t.ozkan,g.agugiaro@tudelft.nl

Keywords: UBEM, Window-to-Wall ratio, Scenario comparison, Cooling demand, Solar irradiance

Abstract

This research investigates how window-to-wall ratio (WWR) granularity affects urban building energy modelling for Positive Energy District planning. We use surface-level WWR data extracted from oblique aerial imagery and compare the resulting simulation outputs against fixed district-average and generic WWR assumptions across two neighbourhoods: Feijenoord and Prinsenland, in Rotterdam, The Netherlands. The analysis uses harmonised building data, adapted to the specific requirements of CitySim and SimStadt simulation tools to evaluate each tool independently against a baseline. Results indicate that fixed WWR assumptions significantly distort simulated energy profiles. The highest effects occur for cooling demand: applying a fixed 20% WWR leads to an increase in the cooling demand by 106.2% in Feijenoord and 127.8% in Prinsenland using CitySim. Such behaviour happens as well in SimStadt, where cooling demand values increase by 67.2% and 68.5%, respectively. Building-level distributions reveal even stronger local deviations, particularly where baseline cooling demand is low. Finally, the findings show that while neighbourhood-average WWR values reduce aggregate error relative to generic assumptions, they still obscure building-specific peaks that are critical for accurate district energy planning.

1. Introduction

As global urbanisation intensifies, cities are projected to house nearly 70% of the world's population by 2050 (UN. Population Division, 2018). At the same time, the increasing adoption of technologies such as electric vehicles and heat pumps directly impacts urban energy demand. Moreover, it places additional pressure on existing utility networks, which are already, in some cases, congested (Li and Jenn, 2024; Hennig et al., 2023). These developments increase the complexity of planning and operating urban energy systems on the district scale.

In this context, Positive Energy Districts (PEDs) have emerged as a relevant planning concept. PEDs are defined as "energy-efficient and energy-flexible urban neighbourhoods or areas of connected buildings and facilities that produce local renewable energy, achieve net-zero greenhouse gas emissions, and actively contribute to overall climate neutrality" (Haindlmaier et al., 2025). Designing and assessing such districts requires accurate simulation of district-level building energy demand, making Urban Building Energy Modelling (UBEM) a key tool in the PED planning process.

However, the effectiveness of UBEM is often faced with data scarcity, particularly when the study area includes many buildings. An important example is the window-to-wall ratio (WWR), defined here as the ratio between the transparent opening area and the total facade area of the thermal envelope of the building. Because detailed window geometry is rarely available on the city scale, WWR is commonly assigned using fixed percentage values based on age, typology, or a single default value. Although this assumption reduces modelling effort, it ignores the spatial heterogeneity of urban facade composition and may affect simulated heating demand, cooling demand, and solar exposure.

This paper analyses the impact of WWR granularity on UBEM outputs for two neighbourhoods in Rotterdam, the Netherlands:

- Feijenoord: an area of 8.5 km^2 with a heterogeneous building stock, including pre-war structures and modern residential blocks; it has 7830 inhabitants and 1648 buildings (van Bijsterveld, 2016a).
- Prinsenland: a 1.8 km^2 area mostly consisting of newer constructions with more homogeneous energy profiles; it has 10,400 inhabitants and 4995 buildings (van Bijsterveld, 2016b).

We present three main contributions. First, we quantify how fixed WWR assumptions alter energy demand estimates compared to facade data extracted from oblique aerial imagery. Second, we evaluate the CitySim and SimStadt engines independently. We avoid direct numerical benchmarking between the two tools, as such a comparison falls outside our scope. Finally, we connect neighbourhood-level effects with building-level distributions and specific extreme cases. We emphasize this connection because aggregate results often mask local deviations that are critical for robust energy-system design.

The analysis is guided by three research questions. The first question asks how much a fixed 20% WWR assumption –a rather typical assumption in such cases– (Veillette et al., 2021) changes annual heating and cooling outputs relative to custom WWR data per surface. The second question asks whether district-average WWR assumptions reduce these changes enough to be useful for aggregate urban energy studies. The third question asks whether neighbourhood-level results are sufficient for PED-oriented interpretation, or whether building-level distributions and selected cases reveal additional risks. These questions are addressed separately per simulation tool.

2. Related Work

Urban building energy modelling bridges detailed building simulation and large-scale building stock analysis. Early reviews describe UBEM as a bottom-up approach that can evaluate neighbourhoods or entire cities by generating many building energy models from geospatial, semantic, and typological datasets (Reinhart and Cerezo Davila, 2016). Later work emphasises that urban models must handle interactions between buildings, district systems, temporal load diversity, and uncertainty in input data (Hong et al., 2020). These requirements are directly relevant for PED studies, where local demand peaks, retrofit priorities, and renewable integration are evaluated on the district scale rather than for isolated buildings.

A recurring limitation in UBEM is the reliance on building archetypes. Archetypes make urban simulation feasible by assigning representative parameters to groups of buildings, but they also simplify the diversity of the building stock. Studies comparing archetype characterisation methods show that classification choices can influence model predictions, especially when the selected archetype variables do not capture important physical differences (Cerezo et al., 2017). Calibration-oriented work further shows that archetypes are useful but should not be treated as exact representations of individual buildings (Kristensen et al., 2018). This research addresses one specific archetype parameter, WWR, and tests whether district-average values can approximate information at the surface level.

WWR is a sensitive design and modelling variable because it affects heat transfer, solar gains, daylight availability, and cooling loads. Parametric-building studies have shown that WWR can be among the influential envelope variables in early-stage energy simulation (Samuelson et al., 2016), while urban-scale studies include WWR as a geometric variable affecting energy consumption (Lila and Lannon, 2017). Occupant behaviour and orientation can moderate the WWR effects, making a single deterministic value problematic for heterogeneous districts (Veillette et al., 2021).

Recent work has started to close this data gap with computer vision. Deep-learning workflows using street-view imagery can estimate WWR at facade, building and district scales and then propagate the resulting values into UBEM simulations (Suppa et al., 2025). Other work explores the generation of more detailed geometry for energy simulation from remote sensing and imagery sources (Kim et al., 2026). These approaches indicate that extracted envelope data can improve physical specificity, but they also raise the question of whether the added detail materially changes district-level energy conclusions.

3. Materials and Methods

3.1 Input Data and Preprocessing

This study uses a dataset provided by the Data Science Department of the Municipality of Rotterdam. The dataset was generated using deep learning algorithms to extract window areas from high-resolution aerial oblique imagery. The extracted WWR data provides a granular representation for building surfaces in the 3D city model. Before simulation, the CityGML LoD2 models are processed to improve geometric and semantic consistency, following the broader workflow described in (Gao et al., 2025). Data preparation includes:

- Geometry repair to correct or mitigate geometric errors in municipal CityGML data using CityDoctor2 (Betz et al., 2026) and FME (Safe Software, 2026).
- Semantic integration to map Dutch basic registry data – BAG in Dutch – (Kadaster, 2026) and Dutch building archetypes (RVO, 2023) to the CityGML buildings to assign and use physics parameters.
- WWR parameterisation to replace standard WWR assumptions with the provided values per surface and two fixed WWR scenarios, one per neighbourhood and a second general for both study areas.

Both engines utilize physical and semantic building parameters, though individual building representations remain tailored to each tool's requirements. We analyze SimStadt (Coors et al., 2021) and CitySim (Kaemco, 2026) independently to assess the internal impact of WWR assumptions, rather than to benchmark the tools against each other.

The study areas offer complementary contexts. Feijenoord's heterogeneous morphology tests whether a district-averaged WWR adequately represents a mixed building stock. Conversely, Prinsenland's highly homogeneous urban fabric evaluates the efficacy of an area-specific average within such environment.

To isolate WWR impacts from other variables, the preprocessing workflow maintains constant physical and semantic parameters, treating WWR as the sole controlled variable. While the respective engines generate distinct model representations and output quantities, this strategy ensures a rigorous, consistent physical baseline for intra-engine scenario comparisons.

3.2 Scenario Design

We define four scenarios based on the available WWR data (table 1). We use the AI scenario (i.e. the set of WWR values obtained from deep learning by the Municipality of Rottedam) as the baseline, as it represents the highest-granularity WWR input available for this research. We apply the AI, 7P, and 20P scenarios to Feijenoord, and the AI, 10P, and 20P scenarios to Prinsenland. The 7P and 10P scenarios represent local averages (corresponding to average WWR rations of 7%, 10% and 20%), allowing us to test whether a neighbourhood-specific archetype can adequately approximate the baseline. Conversely, the 20P scenario represents a generic assumption, of one fixed value without local facade information. We report all results as percentage differences from the baseline:

$$\Delta_{\%} = \frac{X_{\text{scenario}} - X_{\text{AI}}}{X_{\text{AI}}} \times 100 \quad (1)$$

where X is the yearly or monthly value of the analysed metric.

We processed all simulation results using custom Python scripts. CitySim outputs comprise heating, cooling, and solar irradiance, whereas SimStadt outputs heating and cooling. Because SimStadt's *Solar Potential Analysis* showed no

Table 1. Description of the applied WWR scenarios

Scenario	Definition	Neighbourhoods
AI (Extracted)	Surface-specific WWR extracted from oblique aerial imagery using deep learning	Both districts
7P (Mean)	A fixed 7% WWR, representing the mean value for Feijenoord	Feijenoord
10P (Mean)	A fixed 10% WWR, representing the mean value for Prinsenland	Prinsenland
20P (Standard)	A fixed 20% WWR, as traditionally assumed when no data at all are available	Both districts

variation across the three scenarios, which indicates that the computation is done over the LoD2 surface without including additional semantic separation of opaque and galzzing areas. Therefore, we excluded these values from our analysis.

Although annual totals summarise the district energy balance, they can mask compensating errors where building-level overestimations and underestimations cancel each other out. Consequently, we evaluate building-level distributions as a secondary validation to ensure accurate modelling. We consider a scenario robust only if it closely aligns with the baseline on both the neighbourhood and building scales.

3.3 Analysis Levels

The analysis is performed at three levels. The first level is yearly neighbourhood aggregation, which is used to evaluate the overall effect of each WWR scenario. The second level is monthly neighbourhood aggregation, which is used to identify seasonal concentration of the differences. The third level is building-level analysis, which is used to inspect the distribution of percentage changes and to identify selected extreme cases.

We use the AI scenario as our baseline. For all other scenarios, we report the absolute value, absolute delta, and relative percentage difference. Although we prioritise percentages to facilitate comparisons across neighbourhoods and metrics, we include absolute values for selected cases to prevent misinterpretation.

The analysis intentionally avoids direct numerical comparison between CitySim and SimStadt. A direct comparison would require a separate validation design, consistent output definitions, and a discussion of solver and model-structure differences. Instead, we check whether each tool shows sensitivity to WWR assumptions when evaluated against its own baseline. This distinction allows the results to be useful for modelling practice without turning this paper into a software benchmark.

For the building-level analysis, we report percentage differences only when the baseline value is non-zero. When the baseline is zero and a scenario produces a non-zero value, we record the absolute scenario value and describe the case qualitatively rather than reporting an infinite percentage. We use the median and interquartile range as the primary distribution statistics because they are less sensitive to extreme outliers than the mean, which is important when a small number of buildings with near-zero baseline cooling can dominate the average percentage difference.

4. Results

4.1 CitySim Neighbourhood-Level Results

CitySim results show that WWR assumptions affect both heating and cooling on the neighbourhood scale, while solar irradiance varies only slightly. In Feijenoord (figure 1), the 7P scenario reduces heating by 4.70% and cooling by 11.90% relative to the baseline. The 20P scenario increases heating by 25.15% and cooling by 106.23%. In Prinsenland (figure 1), the 10P scenario remains close to the baseline for cooling, with a 1.16% increase, but increases heating by 3.95%. The 20P scenario has a stronger effect, increasing heating by 41.05% and cooling by 127.77%.

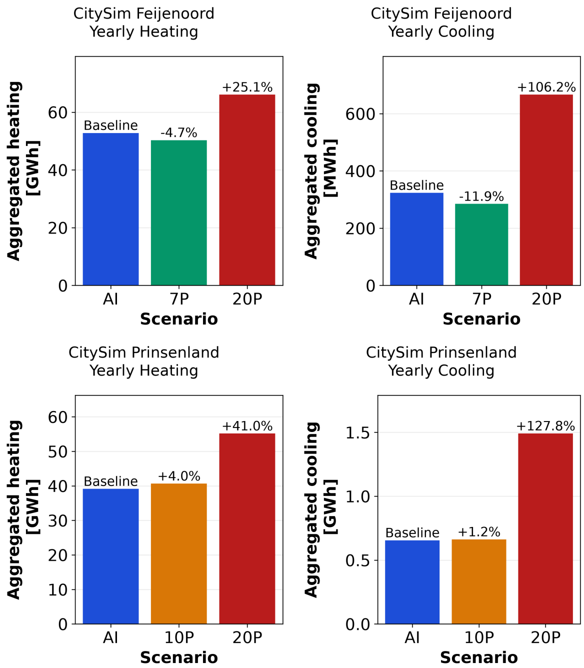

Figure 1. CitySim yearly heating and cooling comparison for Feijenoord and Prinsenland

CitySim results demonstrate that the fixed 20% assumption consistently overestimates annual heating and cooling loads across both neighbourhoods. In Feijenoord, the 7P scenario underestimates both heating and cooling, systematically shifting annual demand downwards rather than merely smoothing individual facade variations. Specifically, it reduces cooling by 11.90%—a deviation potentially acceptable for exploratory studies but significant for detailed energy planning. Conversely, the 20P scenario shifts these outputs upwards by a substantially larger margin. Prinsenland's results highlight that district-average WWRs perform inconsistently across metrics. Although the 10P scenario accurately approximates baseline cooling, it still increases annual heating by 3.95%.

4.2 SimStadt Neighbourhood-Level Results

SimStadt results show a different within-tool balance between heating and cooling. Heating remains comparatively close to the baseline at neighbourhood scale. In Feijenoord (figure 2),

https://doi.org/10.5194/isprs-archives-L-4-W3-2026-83-2026

Table 2. CitySim yearly percentage differences relative to the baseline

Neighbourhood	Scenario	Heating	Cooling	Solar
Feijenoord	7P	-4.70	-11.90	0.35
Feijenoord	20P	25.15	106.23	-0.39
Prinsenland	10P	3.95	1.16	0.32
Prinsenland	20P	41.05	127.77	-0.30

7P reduces heating by 3.12%, while 20P increases heating by 1.35%. In Prinsenland, 10P reduces heating by 5.11%, and 20P remains close to the baseline with a 0.38% decrease. Cooling is more sensitive. In Feijenoord, 7P reduces cooling by 11.53%, while 20P increases cooling by 67.20%. In Prinsenland, 10P reduces cooling by 4.04%, while 20P increases cooling by 68.47%.

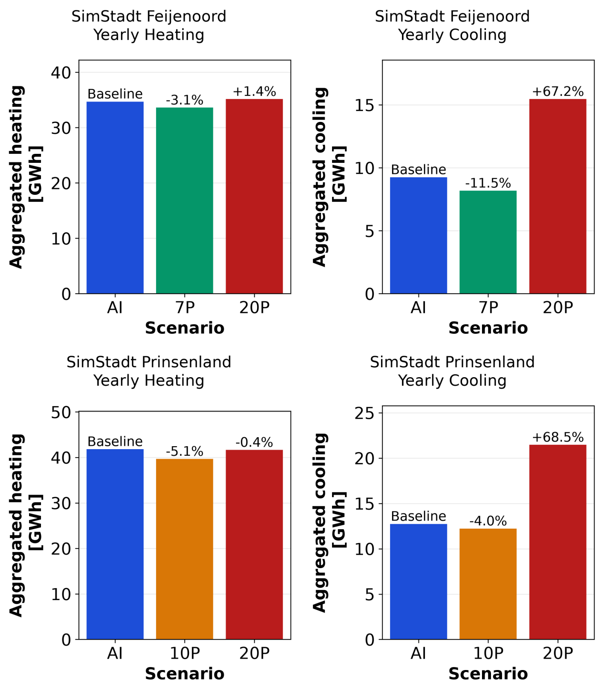

Figure 2. SimStadt yearly heating and cooling comparison for Feijenoord and Prinsenland

The SimStadt results support a different interpretation of the same scenario set. Annual heating is comparatively stable under WWR changes. This suggests that, for the model configuration and input assumptions, heating demand is dominated by other parameters. Cooling, however, reacts strongly to the 20P scenario and its effect is large enough to affect district-scale interpretation.

The district-average scenarios again perform better than the 20P scenario on the neighbourhood scale. Feijenoord 7P reduces annual cooling by 11.53%, while Prinsenland 10P reduces it by 4.04%. These values show that local WWR averages are not perfect substitutes for extracted facade data, but they avoid the much larger cooling increases associated with 20P.

If a planner used the 20P scenario as the basis for cooling infrastructure assumptions, the resulting annual demand estimate would be substantially above the baseline in both neighbourhoods. This supports the central argument that default WWR assumptions should be treated as a source of model uncertainty rather than as neutral input data.

When viewed side by side, the CitySim and SimStadt results show a consistent qualitative pattern: the generic 20P assumption increases cooling demand substantially in both neighbourhoods, while the district-average scenarios stay closer to the baseline. The key message for practice is that both simulation tools are sensitive to WWR assumptions, but the metric most affected—and the size of the effect—depends on the engine and the local building stock configuration.

Table 3. SimStadt yearly percentage differences relative to the baseline

Neighbourhood	Scenario	Heating	Cooling
Feijenoord	7P	-3.12	-11.53
Feijenoord	20P	1.35	67.20
Prinsenland	10P	-5.11	-4.04
Prinsenland	20P	-0.38	68.47

4.3 Monthly Patterns

The monthly outputs show that WWR effects are not evenly distributed across the year. For both engines, the largest cooling deviations occur in spring or summer shoulder months, where baseline cooling can be low, and percentage differences therefore become large. In CitySim, the largest monthly cooling increase for the 20P scenario occurs in April for both neighbourhoods: 284.57% in Feijenoord and 336.74% in Prinsenland. These patterns are shown in figure 3. In SimStadt, the largest monthly cooling increases for 20P also occur in April, reaching 122.82% in Feijenoord and 154.40% in Prinsenland; these monthly differences are shown in figure 3.

For heating, CitySim shows the largest 20P monthly percentage increases in July, with 70.16% in Feijenoord and 131.77% in Prinsenland. These values should be read as seasonal relative changes rather than as dominant annual loads, because summer heating values are low. SimStadt's strongest monthly heating deviations are also found in low-demand periods: the largest 20P decreases occur in June for Feijenoord and July for Prinsenland. The monthly analysis therefore confirms that percentage-based interpretation must be paired with the baseline magnitude. The concentration of large cooling deviations in the spring and early-summer months has practical implications. In these months, solar gains begin to increase while ambient temperatures are still moderate, so buildings with high WWR absorb more solar energy before natural ventilation or shading can fully offset the gains. This means that a fixed high WWR assumption does not only raise annual cooling totals; it can also shift the timing of cooling demand earlier in the year. For district energy planning, this timing matters because the design capacity for cooling systems and the scheduling of demand-response measures are often based on monthly peak profiles rather than annual sums alone.

4.4 CitySim Solar Irradiance

On the neighbourhood scale, the yearly effect of WWR scenario choice on solar irradiance is small compared with the heating and cooling effects. In Feijenoord, 7P increases solar irradiance by 0.35% and 20P reduces it by 0.39%. In Prinsenland, 10P increases solar irradiance by 0.32% and 20P reduces it by 0.30%. The direction is consistent with the geometric logic

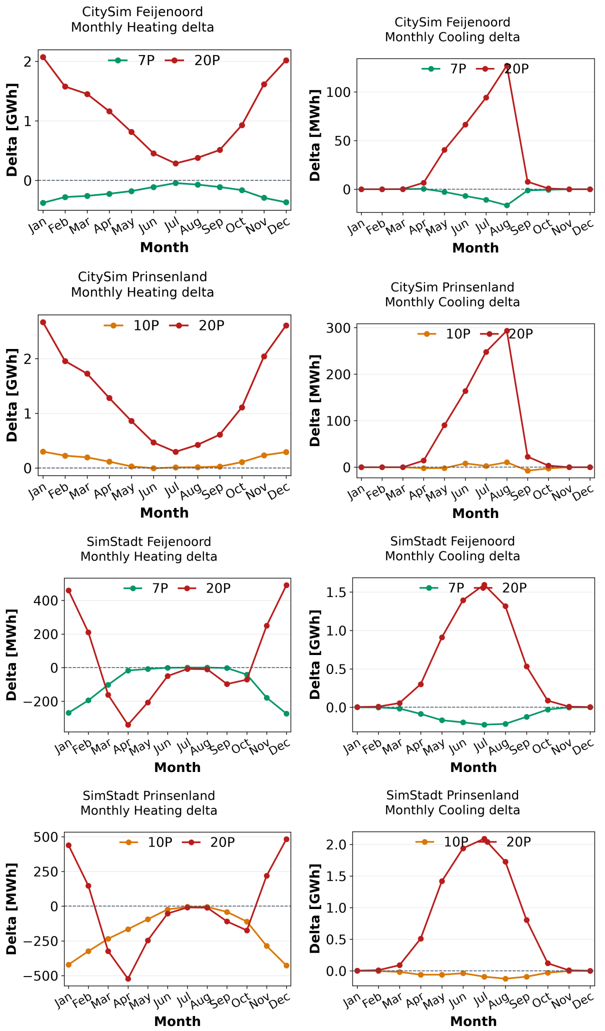

Figure 3. CitySim and SimStadt monthly percentage differences from the baseline for Feijenoord and Prinsenland

of the scenarios: lower WWR leaves more opaque surface available, while higher WWR reduces opaque surface area. However, the aggregate effect remains limited because solar availability is also controlled by roof area, orientation, shading, and urban form.

The small annual percentage differences do not mean that WWR is irrelevant for solar planning. They mean that, on the neighbourhood scale, the change in available opaque facade surface is small relative to the total solar exposure captured by the CitySim outputs. For building-integrated photovoltaics, facade-level geometry matters because the location, orientation, and shading of opaque surfaces determine whether additional area is technically useful. The CitySim solar results therefore add a complementary criterion: a WWR scenario changes demand, but it can also change the surface assumptions used to estimate local generation potential.

4.5 Building-Level Distributions

The building-level distribution summaries show that neighbourhood averages hide large local variation. The scope of figure 4 is to visualise the spread of building-level cooling percentage differences relative to the baseline for each WWR scenario. These figures show how widely individual buildings deviate from the baseline within each simulation tool. In CitySim, the median cooling difference for 20P is 202.60% in Feijenoord and 232.17% in Prinsenland, with extreme maximum values caused by buildings whose baseline cooling demand is close to zero. Heating is less extreme than cooling but still relevant: the median 20P heating difference is 38.13% in Feijenoord and 55.75% in Prinsenland. Solar irradiance has much narrower building-level distributions, with median differences below one percent in both neighbourhoods.

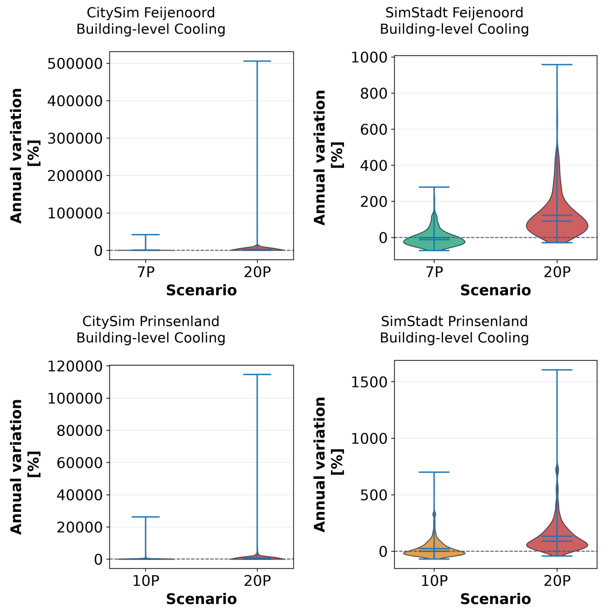

Figure 4. Building-level cooling percentage distributions in Feijenoord and Prinsenland, analysed separately for CitySim and SimStadt

In SimStadt, the median 20P cooling difference is 89.46% in Feijenoord and 89.10% in Prinsenland. Heating distributions are narrower: the median 20P heating difference is 0.93% in Feijenoord and 2.29% in Prinsenland. The district-average scenarios are closer to the baseline at aggregate scale, but their building-level distributions still include strong positive and negative deviations. This is visible in the Prinsenland distribution in figure 4: the 10P cooling distribution in SimStadt has a median close to zero, but the range extends from -69.32% to 700.00%.

These distributions are important because they show the difference between aggregate representativeness and local accuracy. A district-average WWR can be representative in the statistical sense that the total neighbourhood value remains close to the baseline outputs. It can still be locally incorrect for individual buildings whose facade composition differs from the average. The Feijenoord and Prinsenland distributions in figure 4 are therefore used to assess local variability, not only neighbourhood totals. This distinction matters whenever modelling results are used to prioritise interventions, identify critical buildings, or evaluate local grid reinforcement.

The interquartile ranges provide a more stable view than the extreme values alone. For CitySim Feijenoord cooling under 20P, the middle half of buildings ranges from 48.24% to 534.04%. For CitySim Prinsenland cooling under 20P, the middle half ranges from 101.46% to 519.44%. In SimStadt, the corresponding 20P cooling interquartile ranges are 42.66% to 156.68% in Feijenoord and 43.57% to 170.25% in Prinsenland.

These ranges, visualised in figure 4, show that the cooling sensitivity is not limited to a few extreme outliers. The extremes are useful diagnostic cases, but the broader distributions show that WWR simplification affects many buildings.

4.6 Selected Extreme Building Cases

The selected building cases illustrate why both aggregate and disaggregated analysis are needed. In CitySim, the largest cooling increases are dominated by very small baseline values. In Feijenoord, building with id `BAG_0599100100010065` varies from 0.017 to 85.960 in the 20P scenario, which corresponds to 505,547.03%. In Prinsenland, building with id `BAG_0599100000345093` varies from 0.543 to 623.140 in the 20P scenario, corresponding to 114,658.75%. The monthly behaviour of these cases are shown in figure 5. These percentages are mathematically correct, but their interpretation depends on the absolute baseline magnitude. When checking SimStadt results, in Feijenoord, `BAG_0599100000034717` increases by 957.24% under 20P. In Prinsenland, `BAG_0599100000755297` increases by 1604.76% under 20P. These maximum-increase cases are shown in figure 5.

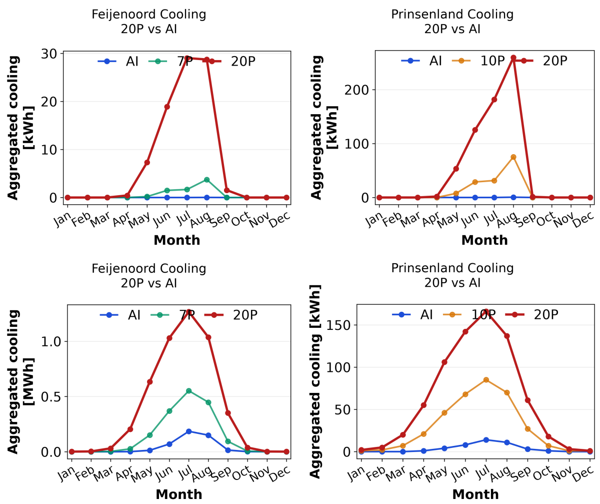

Figure 5. Monthly cooling demand of buildings in CitySim and SimStadt for the maximum-difference cases under the 20P scenario for Feijenoord and Prinsenland

Several buildings are present in both simulation outputs as huge difference cases for cooling demand. These shared cases should be interpreted qualitatively, not as numerical comparisons between tools. For example, `BAG_0599100000755297` in Prinsenland shows a strong cooling increase under 20P in both outputs, indicating that the building is sensitive to the generic WWR assumption.

Table 4. Examples of selected cooling extreme cases reported relative to the baseline

Engine	Area	Case	Difference (%)
CitySim	Feijenoord	20P max diff.	505547.03
CitySim	Prinsenland	20P max diff.	114658.75
SimStadt	Feijenoord	20P max diff.	957.24
SimStadt	Prinsenland	20P max diff.	1604.76

The selected cases also show that the same WWR scenario can increase one metric while reducing another, depending on baseline conditions and model response. Such cases are relevant for PED planning because interventions are rarely evaluated against a single metric. A facade assumption that reduces one demand component may increase another, or it may change demand while leaving solar availability almost unchanged.

Table 4 is intentionally limited to examples rather than a complete ranking. The table illustrates why selected cases require careful explanation: very large percentages can occur when the cooling baseline is nearly zero, so relative values must be interpreted together with absolute baseline and scenario values.

The presence of the same building identifier in both the CitySim and SimStadt extreme-case sets—such as `BAG_0599100000755297` in Prinsenland—is particularly informative. It indicates that certain buildings are sensitive to the generic WWR assumption regardless of the specific solver or model structure. This qualitative agreement strengthens the argument that the 20% assumption can create locally problematic demand estimates. At the same time, the absolute values and monthly patterns differ between engines, which reinforces why we do not numerically benchmark one tool against the other.

5. Discussion

The results indicate that WWR values affects urban energy simulation in two related but distinct ways. On the neighbourhood scale, the fixed 20% assumption consistently changes cooling demand more strongly than district-average scenarios. This supports the use of local WWR statistics when extracted facade data are unavailable. A 7% value for Feijenoord and a 10% value for Prinsenland reduce the aggregate deviation from baseline, especially for cooling. However, they do not eliminate building-level deviations.

The second effect concerns spatial heterogeneity. Building-level distributions and selected cases show that a scenario can appear acceptable on the neighbourhood scale while still creating large local deviations. This is important for PED planning because district energy systems are affected by peak loads, local grid constraints, and the spatial clustering of demand. A neighbourhood-average WWR may be sufficient for first-order annual energy accounting, but it is less reliable when the objective is to size local cooling systems, identify building-level retrofit priorities, or evaluate flexibility at specific locations.

The results also show why percentage differences must be handled carefully. Very large building-level cooling percentages often occur when the baseline is close to zero. These cases are relevant because a small baseline load can become non-negligible under a simplified WWR assumption, but the percentage alone exaggerates the apparent physical magnitude.

For CitySim, solar irradiance adds a complementary interpretation. The yearly solar irradiance changes are small on the neighbourhood scale, but they follow the expected direction of the WWR scenarios. Lower WWR leaves more opaque area and slightly increases the available irradiance on opaque surfaces, while higher WWR reduces it. The limited magnitude suggests that, in these districts, solar irradiance availability is controlled by a wider set of geometric factors.

Taken together, the neighbourhood-level, monthly, and building-level results suggest that the risk of WWR simplification is not uniform across modelling tasks. The largest risks appear when the goal is to estimate cooling demand, identify local peaks, or support building-level decision-making. In these cases, even district-average WWRs can produce noticeable deviations. For first-order annual heating estimates, by contrast, the tested WWR assumptions have a smaller aggregate effect, at least in the SimStadt configuration used here. This asymmetry implies that modellers should choose their WWR input strategy according to the specific question being asked, rather than applying a single default value as a universal simplification.

The findings also underline the difference between sensitivity and accuracy. This paper quantifies sensitivity to WWR assumptions, treating the AI scenario as the most detailed available baseline. It does not claim that the AI scenario is an observational ground truth for real energy consumption. A separate validation against metered data would be needed to assess absolute accuracy. However, the large sensitivity to WWR assumptions is itself a meaningful result: even if the absolute values were calibrated, the spread introduced by fixed WWR assumptions would remain a source of structural uncertainty in the model outputs.

5.1 Implications for Data Acquisition

Extracted facade data are necessary depending on the goal. A district-average WWR works for rough estimates of annual neighborhood demand, especially in homogeneous areas like Prinsenland. However, for estimating cooling peaks, evaluating building-level retrofits, or planning local grids, extracted WWR is essential because building-level variations remain significant even when aggregate numbers look similar.

The fixed 20% WWR assumption performed worst across all test scenarios, causing large cooling deviations in both neighborhoods and modelling chains. While not inherently wrong everywhere, a generic 20% assumption fails to represent these Rotterdam districts accurately.

These results can guide data collection priorities. Extracting facade-level WWR from imagery requires significant time and resources, which may not be justified for studies focused only on annual energy balance, where a locally calibrated average is often sufficient. However, for analyses of cooling peaks, retrofit prioritisation, or grid reinforcement, facade-level detail becomes more valuable. A practical approach is to use district-average WWRs for initial screening and extract facade-level data only for priority buildings or critical areas, balancing data costs with the required level of detail.

5.2 Implications for PED Planning

PED planning requires simultaneous attention to demand reduction, renewable generation, and energy flexibility, all of which are indirectly affected by WWR assumptions. WWR influences both sides of the energy equation—heating and cooling loads on the demand side, and available opaque surface for local solar generation on the supply side. Because district energy targets are implemented through building-level decisions (like retrofits, heat pumps, and solar integration), neighbourhood-average WWR values can be misleading. A model might predict the total district energy accurately while misallocating demand among individual buildings, distorting project prioritization. Thus, district-level accuracy alone shouldn't be the sole validation benchmark for urban energy models.

WWR creates a direct trade-off between energy demand and generation: higher WWR increases cooling loads while reducing space for facade-integrated photovoltaics. While this trade-off was minor in the studied districts, it becomes critical in dense urban areas with limited facade space and high cooling demands. PED assessments must treat WWR as a parameter that shapes both energy supply and demand. Given that WWR assumptions can swing cooling estimates by over 60% to 100% depending on the modelling engine, PED workflows must include sensitivity testing. Reporting this uncertainty alongside main results ensures decision-makers know whether a projected energy performance is truly robust or merely an artefact of simplified modelling.

6. Limitations

This study relies solely on existing processed scenario outputs without re-running simulations or recalibrating parameters. Future research should test if these patterns hold under different schedules, weather data, infiltration rates, glazing types, or cooling set-points. In addition, simulated demand was not validated against measured energy consumption. The AI scenario serves as a detailed WWR baseline rather than ground truth, as our focus is WWR sensitivity, not absolute model accuracy. Full validation would require actual consumption data, occupancy details, and calibration protocols.

Although more granular than fixed scenarios, the extracted facade data may still contain detection errors, occlusions, or semantic misalignments. While treated here as the best available comparison point, future work should quantify and propagate WWR extraction uncertainty directly into energy simulations. Direct numerical comparisons between CitySim and SimStadt were deliberately avoided due to their differing modelling structures and input data. A fair absolute comparison requires a dedicated benchmarking protocol; thus, this paper focuses on internal scenario sensitivity within each tool.

7. Conclusion

In this research, we examined how extracted and fixed WWR assumptions affect UBEM results in two Rotterdam neighbourhoods. The findings show that the 20P generic assumption can substantially alter cooling demand relative to extracted facade data. District-average WWR values reduce aggregate deviations, but they still mask building-level variability and selected extreme responses.

The analysis supports three practical conclusions. First, extracted WWR is valuable when the planning question depends on local peaks, building-level prioritisation, or spatially explicit cooling demand. Second, neighbourhood-specific WWR averages are a defensible fallback for aggregate analysis when detailed facade data are unavailable. Third, simulation-engine outputs should be interpreted within each modelling chain; qualitative tendencies can be discussed, but direct numerical benchmarking between CitySim and SimStadt is outside the scope of this research.

Future work should connect the selected-building deviations to facade morphology, building use, and local shading conditions.

It should also test whether the same patterns hold in districts with different construction periods, urban densities, and cooling adoption assumptions.

8. Acknowledgements

The work presented in this paper was developed within the European projects ”DigiTwins4PEDs” (Utilisation of urban digital twins to co-create flexible positive energy systems for districts), DUT call 2022, and ”FlexPED” (Energy Data and Flexibility Management for Positive Energy Districts), DUT call 2024. Both projects have been funded in the Netherlands by the Dutch Research Council (NWO) under the Driving Urban Transitions Partnership, which is co-funded by the European Commission.

The authors would like to thank the municipality of Rotterdam for their support as well as providing us access to the data used for this research. Specifically to Milad Mahour, Martijn de Bakker, Thomas M. Kocken, Matthaeus M. Hilpold, Roel Westrik from Research and Business Intelligence Department as well as Marie-Emilie Ingen Housz Heat and Cooling Advisor for her commitment and facilitating the connection with the municipality.

References

Betz, M., Riegel, L., Coors, V., 2026. CityDoctor / CityDoctor2 · GitLab. url: https://transfer.hft-stuttgart.de/gitlab/citydoctor/citydoctor2.

Cerezo, C., Sokol, J., AlKhaled, S., Reinhart, C., Al-Mumin, A., Hajiah, A., 2017. Comparison of four building archetype characterization methods in urban building energy modeling (UBEM): A residential case study in Kuwait City. *Energy and Buildings*, 154, 321–334. https://doi.org/10.1016/j.enbuild.2017.08.029. doi: 10.1016/j.enbuild.2017.08.029.

Coors, V., Rodrigues, P., Betz, M., Schneider, S., Weiler, V., Duminil, E., Schröter, B., Klöber, A., Holweg, D., Brüggemann, T., Bohn, K., Rein, M., Goll, L., Balbach, B., Spath, F., Gärtner, B., 2021. EnEff:Stadt – SimStadt 2.0. Technical Report 03ET1459, Hochschule für Technik Stuttgart (HFT).

Gao, W., Agugiaro, G., León-Sánchez, C., 2025. Data-driven energy simulations to evaluate positive energy district potential in Rotterdam. *ISPRS Annals of the Photogrammetry, Remote Sensing and Spatial Information Sciences*, X-4/W7-2025, Tokyo, Japan, 25–32. doi: 10.5194/isprs-annals-X-4-W7-2025-25-2025.

Haindlmaier, G., Reiter, C., Hinterberger, R., 2025. PED Framework 3.0. Policy Guide, Driving Urban Transitions (DUT) Partnership for SET Plan IWG 3.2.

Hennig, R. J., De Vries, L. J., Tindemans, S. H., 2023. Congestion management in electricity distribution networks: Smart tariffs, local markets and direct control. *Utilities Policy*, 85, 101660. doi: 10.1016/j.jup.2023.101660.

Hong, T., Chen, Y., Luo, X., Luo, N., Lee, S. H., 2020. Ten questions on urban building energy modeling. *Building and Environment*, 168, 106508. https://doi.org/10.1016/j.buildenv.2019.106508. doi: 10.1016/j.buildenv.2019.106508.

Kadaster, 2026. Dataset: Basisregistratie Adressen en Gebouwen (BAG).

Kaemco, 2026. kaemco | Download... Available at: http://kaemco.ch/download.php (accessed 9 April 2026).

Kim, Y., Yoon, J., Lee, S., Shin, M., 2026. Automated building geometry modeling using UAV imagery and deep learning for energy simulation. *Journal of Building Engineering*, 117, 114949. https://doi.org/10.1016/j.jobe.2025.114949. doi: 10.1016/j.jobe.2025.114949.

Kristensen, M. H., Hedegaard, R. E., Petersen, S., 2018. Hierarchical calibration of archetypes for urban building energy modeling. *Energy and Buildings*, 175, 219–234. https://doi.org/10.1016/j.enbuild.2018.07.030. doi: 10.1016/j.enbuild.2018.07.030.

Li, Y., Jenn, A., 2024. Impact of electric vehicle charging demand on power distribution grid congestion. *Proceedings of the National Academy of Sciences*, 121(18), e2317599121. doi: 10.1073/pnas.2317599121.

Lila, A. M., Lannon, S., 2017. A parametric sensitivity analysis of the impact of built environment geometrical variables on building energy consumption. *Proceedings of Building Simulation 2017*, 2, 2092.

Reinhart, C. F., Cerezo Davila, C., 2016. Urban building energy modeling – A review of a nascent field. *Building and Environment*, 97, 196–202. https://doi.org/10.1016/j.buildenv.2015.12.001. doi: 10.1016/j.buildenv.2015.12.001.

RVO, 2023. Voorbeeldwoningen 2022 bestaande bouw. Verantwoordingsrapportage W/E 30751 | RVO EPBD 190035, Rijksdienst voor Ondernemend Nederland, Eindhoven. Rijksdienst voor Ondernemend Nederland.

Safe Software, 2026. FME©. Available at: https://fme.safe.com/ (accessed 9 April 2026).

Samuelson, H., Claussnitzer, S., Goyal, A., Chen, Y., Romo-Castillo, A., 2016. Parametric energy simulation in early design: High-rise residential buildings in urban contexts. *Building and Environment*, 101, 19–31. doi: https://doi.org/10.1016/j.buildenv.2016.02.018.

Suppa, A. R., Aliberti, A., Bottero, M. C., Corrado, V., 2025. Detecting window-to-wall ratio for urban-scale building simulations using deep learning with street view imagery and an automatic classification algorithm. *Building Simulation*, 18, 2175–2199. https://doi.org/10.1007/s12273-025-1301-3. doi: 10.1007/s12273-025-1301-3.

UN. Population Division, 2018. *The World 's Cities in 2018*. United Nations, New York, USA.

van Bijsterveld, W., 2016a. Buurt Feijenoord (gemeente Rotterdam) in cijfers en grafieken.

van Bijsterveld, W., 2016b. Buurt Prinsenland (gemeente Rotterdam) in cijfers en grafieken.

Veillette, D., Rouleau, J., Gosselin, L., 2021. Impact of Window-to-Wall Ratio on Heating Demand and Thermal Comfort When Considering a Variety of Occupant Behavior Profiles. *Frontiers in Sustainable Cities*, 3, 700794. https://doi.org/10.3389/frsc.2021.700794. doi: 10.3389/frsc.2021.700794.